contact@namiq.io · namiq.io

NamiQ

SECURITY ADVISORY

N-able MailAssure exploited

CVE-2026-134624 · on CISA's exploited list since 23 November 2026

Analysis Report

Operators of N-able MailAssure on any build before 20261121.1943 must act by 26 November 2026. If no external or untrusted users hold API credentials the exposure is reduced, but update to build …

PREPARED BY	NamiQ	DATE	25 November 2026
FACTS AS OF	25 November 2026	VERSION	1.0
SEVERITY	**High: CVSS 3.1 8.7**	STATUS	On CISA's exploited list; due 26 November 2026

NamiQ © 2026 NamiQ. All rights reserved. This document is the property of NamiQ. It may be shared unaltered, with attribution.
Companion article: namiq.io/insights/n-able-mailassure-cve-2026-134624

The bottom line

> *Operators of N-able MailAssure on any build before 20261121.1943 must act by 26 November 2026. If no external or untrusted users hold API credentials the exposure is reduced, but update to build 20261121.1943 or later immediately.*

At a glance

Fact	Detail
Product	N-able MailAssure (CVE-2026-134624)
Severity	High: 8.7 (CVSS 3.1)
CISA	On its exploited list since 23 November 2026. Deadline for US federal agencies: 26 November 2026.
Chance of exploitation (EPSS)	0.02, the estimated chance in the next 30 days, as at 25 November 2026
Ransomware use	Not recorded by CISA. That is not the same as none.
Fixed in	The record names all builds before 20261121.1943 as affected; 20261121.1943 is the first unaffected build. No vendor advisory page was available to confirm further detail.
Who looks first	Anyone running N-able MailAssure (also sold as SpamExperts) on a build earlier than 20261121.1943, especially where API credentials are granted to users outside the core team or where the server can reach sensitive internal services.

Does this apply to you?

Answer in order. Stop at the first no.

Question	If no	If yes
Do you run N-able MailAssure (also known as SpamExperts)?	Nothing to do for this CVE.	Check your build number.
Is your MailAssure build number earlier than 20261121.1943?	No action needed for this CVE.	Update to build 20261121.1943 or later.
Do any users outside your own team hold MailAssure API credentials?	Risk is limited to internal users; still patch.	Treat as urgent; API access is the required condition.

What it is

An authenticated API user can direct the MailAssure server to open connections to any internal host and port it can reach, check which are reachable, and send crafted data to those services. The server becomes a proxy for probing the internal network.

Why it matters

The record carries a CVSS 4.0 score of 9.3. SSRF on a mail-filtering server matters because such servers sit inside the perimeter and can often reach internal services that are not exposed externally. An attacker with a valid API account could map which internal hosts and ports the server can contact, then send crafted data to services that may trust traffic from the mail server. The authentication requirement raises the bar: a credential is needed before anything else. Our reading is that the key question is how tightly API credentials are controlled and what internal services the MailAssure server can reach.

Things people will get wrong

- No outside party has our API credentials, so we are not exposed. The authentication requirement reduces risk but does not remove it. Any compromised internal account with API access is sufficient; the record makes no distinction on the type of API user.

The wording includes conditions the score may not show: the attacker must be logged in. Read the record before you decide it does not apply to you.

How we would triage it

This is our judgement.

Step 1: Find the MailAssure build number you run

- Check the MailAssure admin interface or package information for the installed build number and compare it with 20261121.1943.
- If the build number is not visible in the UI, ask your hosting provider or N-able support for the installed version.

Step 2: Audit who holds API credentials

- List all users with API access in MailAssure and confirm whether any are outside direct team control.
- Revoke or restrict credentials for any account that does not need API access while the patch is applied.

Step 3: Apply the update

- Update MailAssure to build 20261121.1943 or later following your normal change process.
- No vendor advisory was available; confirm the update path with N-able or your hosting provider.

Step 4: Check for signs of misuse if the system was exposed

- Review MailAssure logs for LDAP synchronisation connection check requests directed at internal IP addresses or ports not normally used for LDAP.
- Review network or firewall logs for outbound connections from the MailAssure server to internal hosts it would not normally contact.

To close it, read the version from the product itself, not from an old inventory, and save the output with the date. If you cannot update today, write down who owns it, the new date, and what limits exposure until then. CISA's wording adds that you should stop using the product if no mitigation is available.

Send this to your IT provider

Copy it, change the names and dates, and send it. It is a draft for you to check, not a template tested with any provider.

> *We have been made aware of CVE-2026-134624, a server-side request forgery flaw in N-able MailAssure (SpamExperts) affecting all builds before 20261121.1943. CISA has set a remediation deadline of 26 November 2026. Could you please confirm: (1) whether we run MailAssure and on which build; (2) whether that build is earlier than 20261121.1943; and, if so, (3) when the update to build 20261121.1943 or later can be applied. We would like written confirmation of the installed version and the outcome of any update by 26 November 2026.*

Signs it was used

An update closes the hole. It does not tell you whether someone used it first. CISA's required action points to its forensic triage requirements (BOD 26-04). Look for:

- Check network or firewall logs for outbound connections from the MailAssure server to internal hosts that it would not normally contact.

The weakness behind it

The record classes this as CWE-918 (Server-Side Request Forgery).

Likely ATT&CK mapping

MITRE ATT&CK is a public catalogue of the steps attackers take, from getting in to causing damage. No source has mapped this flaw to it yet, so the steps below are our prediction from what the record says the flaw allows. They are not MITRE's, the vendor's or CISA's, and an attacker may do other things.

Steps an attacker would likely take with this flaw

The step, in plain words	Stage of an attack	Why we think so
Remote System Discovery (T1018): Finding the other machines on the network.	Discovery: looking around	The record says the flaw lets an authenticated user determine which internal hosts and ports are reachable, which would typically be used for internal network discovery.
Exploitation of Remote Services (T1210): Exploiting a flaw in a service on another machine to move to it.	Lateral Movement: moving to other machines	The description says crafted data can be sent to internal services via the SSRF, which would typically be used to exploit services that trust traffic from the mail server.

How to use this: ask your IT provider which control or log would show each step, for example whether sign-in attempts from the internet are recorded. If nobody can answer, that is the gap to close first.

MITRE ATT&CK® is a registered trademark of The MITRE Corporation. Technique names are MITRE's.

The sources, in their words

CISA's catalogue entry:

> *N-able MailAssure (SpamExperts) contains a server-side request forgery vulnerability in the LDAP synchronization connection check that allows an authenticated API user to connect to arbitrary internal hosts and ports, determine which are reachable, and send crafted data to internal services.*

The CVE record, published by N-able on 22 November 2026:

> *N-able MailAssure (SpamExperts) before 20261121.1943 does not restrict the destination of the LDAP synchronisation connection check performed when a relay is configured. An authenticated API user can cause the server to connect to arbitrary internal hosts and ports, determine which are reachable, and send crafted data to internal services.*

Versions of N-able MailAssure the record lists as affected:

- all versions before 20261121.1943

The affected ranges end before 20261121.1943. Those are the first versions the record does not list as affected; confirm them in the vendor's advisory.

CISA's required action:

> *Apply mitigations in accordance with vendor instructions.*

Check your own project

If you build software, you can compare what a project declares with this CVE from a terminal. The tool reads requirements files, package.json, Maven, Gradle, Cargo, CMake and Dockerfile files, runs nothing, and only the CVE id leaves your machine. Naming the CVEs needs no key.

```
lattice-scan project . --cve CVE-2026-134624
```

Each row says in-range, outside-range, undetermined or lead, with the file and line. That is evidence about what a manifest declares, not about what runs, and a person decides. The verdicts, the CSV columns and a pipeline recipe are in the guide.

Check the servers you run

The same id can be checked against a host. Run this on the machine itself, not on the laptop that deploys to it: it reads installed packages (dpkg, rpm or apk), snaps, the hardware (maker, model, firmware), the running container images with their operating system and packages, and the kernel version, compares them on the spot, and only the CVE ids leave the machine; the inventory stays on it, in ~/.lattice-scan. Naming the CVEs needs no key. Do it on each server after a deployment, and again when a new entry lands.

```
lattice-scan --cve CVE-2026-134624
```

A match is a candidate for a person to confirm, not a finding. A distribution can fix a CVE under an unchanged version number, so a match is a candidate; on Debian-family hosts the report also shows whether the package's own changelog names the CVE. By default it opens the images your containers are running and reads their operating system and package list; a deep scan (--deep, a few minutes) opens every local image and the Python, Node, Java and Ruby packages inside them; for what an image declares, use the project check on its Dockerfile.

Guide: namiq.io/lattice/guides/choose-your-channel. MCP server (io.namiq/lattice in the MCP registry): namiq.io/lattice. Package: pypi.org/project/lattice-scan, or namiq.io/apt.

What we do not know yet

- No vendor advisory page was available; the fixed build number comes from the CVE record alone and has not been independently confirmed by a vendor source.
- The record does not say whether on-premises and cloud-hosted customers must both apply an update, or whether cloud customers are patched automatically.
- It is not stated which internal services or protocols are reachable via the LDAP synchronisation endpoint in a typical deployment.
- The record does not describe whether any workaround short of patching is available.
- We have not tested this against any environment. Versions and scores are as the sources state them and as at 25 November 2026; the vendor's advisory wins where it differs.
- The sections headed "The bottom line", "Does this apply to you?", "What it is", "Why it matters", "Exposure by situation", "Things people will get wrong", "Likely ATT&CK mapping", "How we would triage it" and "Send this to your IT provider" are our interpretation of the sources, not a statement by CISA or the vendor.

Appendix 1. Severity decoded

CVE-2026-134624: score 9.3, `CVSS:4.0/AV:N/AC:L/AT:N/PR:H/UI:N/VC:H/VI:H/VA:N/SC:H/SI:H/SA:N` As the CVE record gives it. The record's own score can differ from the base score on the NVD record.

Metric	Value	What it means here
Attack vector	Network	Reachable over a network, with no need to be on the same segment.
Attack complexity	Low	No special conditions beyond the flaw itself.
Attack requirements	None	Works against a default deployment.
Privileges required	High	An administrator account is needed.
User interaction	None	No one has to do anything.
Vulnerable system: confidentiality	High	Total loss of confidentiality.
Vulnerable system: integrity	High	Total loss of integrity.
Vulnerable system: availability	None	No effect on availability.
Subsequent systems: confidentiality	High	Total loss of confidentiality.
Subsequent systems: integrity	High	Total loss of integrity.
Subsequent systems: availability	None	No effect on availability.

Appendix 2. Timeline

Date	Event
22 November 2026	CVE record published by N-able.
23 November 2026	CISA added it to its list of exploited vulnerabilities.
26 November 2026	CISA's deadline for US federal agencies.

Appendix 3. Exposure by situation

These priorities are our judgement from the attacker prerequisites in the sources. They are not vendor ratings.

Your situation	Priority	Timeframe

MailAssure before 20261121.1943 with API credentials issued to external or partner users	High	Within CISA's window — by 26 November 2026
MailAssure before 20261121.1943, server can reach internal management interfaces or databases	High	Within CISA's window — by 26 November 2026
MailAssure before 20261121.1943, API access restricted to internal staff only	Medium	This week
MailAssure at build 20261121.1943 or later	None	No action for this CVE

Appendix 4. Definition of done

- Every N-able MailAssure you run is listed, with its version.
- Each is outside the affected list, or the vendor's mitigation is in place.
- The version output or a screenshot is saved with the date.
- Your projects' declared dependencies were checked against this CVE, and every row has a written reason.
- Logs for the weeks before the update were searched, and anything odd was written down.
- Anything not done has an owner and a date.

Sources

- CISA Known Exploited Vulnerabilities catalogue: https://www.cisa.gov/known-exploited-vulnerabilities-catalog
- CVE record for CVE-2026-134624: https://www.cve.org/CVERecord?id=CVE-2026-134624
- FIRST: Exploit Prediction Scoring System: https://www.first.org/epss/
- Reference in the CVE record: me.n-able.com: https://me.n-able.com/s/security-advisory/aArVy0000002mrlKAA